1 Introduction

Linux is an operating system that runs on IBM-PC compatible machines with Intel-386/486/Pentium or equivalent processors. It was written by Linus Torvalds[1], a graduate student at the Helsinki University of Technology in Finland. Since it was written from scratch, it does not contain any proprietary code. Torvalds holds the copyright on Linux but permits the source code to be freely distributed by anyone. In addition to Torvalds, an enthusiastic world-wide group of volunteers has collaborated in developing this operating system through the Internet. Many of the powerful compilers and utilities developed by the Free Software Foundation under its 'GNU' project can run on Linux. The X Window System from MIT, which is the defacto standard for windowing under Unix, is also available for Linux in an Intel-386/486 version called XFree86 (again because of the selfless efforts of numerous volunteers).

Linux is thus a fully fledged operating system that provides all the capabilities normally associated with commercial Unix systems. It has excellent compilers for C, C++, Pascal, Modula-2 and -3, Oberon, Smalltalk, etc. Fortran is supported via a very reliable Fortran-to-C translator. Editors like `emacs` and `vi` are available in several versions. Password security, file protection, multiple logins, virtual memory and multitasking are fully supported. Networking for Linux has developed to the stage that large networks can be set up to support remote logins, remote shells, electronic mail, talk, etc. A version of the Network File System (NFS) has been developed for Linux. This permits file systems to be shared among several machines, thus reducing hard disk usage and minimizing system administration overhead. The TEX/LATEX text processing system has been ported to Linux along with previewing facilities (`ghostview` & `xdvi`) and drawing packages (`xfig` & `idraw`).

All capabilities described above are available free and without any licensing fee. Furthermore, Linux runs on hardware that is common, cheap and can be purchased anywhere in the world. It thus provides a valuable resource for the Unix professional, who can use on his personal home machine or on a laptop, a system whose interface is the same as that of the workstation in his office. For cost-conscious educational institutions, Linux permits large

[1]Linux is, obviously, named after its author. It is pronounced 'Lee-nucks' in Finland and 'Lie-nucks' in other parts of the world.

laboratories of X Window Systems to be set up on low-cost hardware. Students can afford the full capabilities of Unix on their personal machines, thus reducing load on their institutions' systems. Computer Science and Engineering students can study, modify and recompile the kernel of the operating system, the source code for which is freely available.

In many of the developing world's universities, which cannot afford commercial workstations, Linux may well be the only feasible way to introduce Unix and the X Window System to large numbers of students. For computer enthusiasts all over the world, Linux provides a low risk, low cost path to the exploration of the world of computing beyond DOS and MSWindows. Finally, organizations that have already invested in networks of proprietary Unix/X machines can add Linux based PCs to augment their systems at very low cost.

Linux has been under development for only a few years. As is to be expected, the initial versions were not stable and were of interest only to die-hard 'hackers'. However, over the period 1994-95, Linux has matured to the point where it can be used for serious, routine work. Information about Linux has, by and large, been restricted to discussions on the Internet. There has been little or no coverage in the 'mainstream' technical literature. What little has been written has been unnecessarily discouraging. Obviously, as Linux is not a commercial product, no significant profit is made from its sale. Thus the few advertisements that one sees are from companies that make a modest income from selling diskette sets or CDROMs and are usually in obscure corners of trade magazines.

This paper attempts to remedy the situation by describing how Linux can be obtained, installed and utilized. It is based on my experiences with Linux over the past two years. As I had no prior experience with Unix System Administration but was able to install nearly two dozen Linuxes without any serious problems, I feel this should encourage others to experiment with, and utilize, this shared world-wide resource. This is especially important in view of the plummeting costs of IBM-PC compatible machines. Such machines are easy to purchase and maintain in almost any corner of the world. In 1995, a machine with an Intel 486/DX2-66 processor, 14 inch high definition color monitor, video accelerator card, 16 MB RAM, 500 MB hard disk, one floppy drive, Ethernet card and mouse can be purchased for under US$2000. Such a machine can run Linux and X with superb performance and, other

than its smaller screen[2], seriously challenges traditional workstations. By substituting a slower processor, reducing RAM to 8 MB, using a black and white monitor and 320 MB hard disks, a machine capable of running X with very acceptable performance for teaching purposes can be put together for under US$1400. These are figures no college or university administrator can ignore.

Several books on Linux have been published very recently. Welsh & Kaufman [10] and Strobel & Uhl [11] provide comprehensive details about obtaining, installing and running Linux. Kirsch [6] deals with the important topic of network administration. In addition several fine books and manuals are available in electronic form on the Internet (see Section 2).

2 Obtaining Information about Linux

As befits an operating system that was largely developed over the Internet, Linux, as well as vast amounts of information *about* Linux, are accessible on the Internet. The three main sites are shown below.

Site	Directory	Subdirectory for		
		HOWTOs	FAQ	Books
`ftp.funet.fi`	`/pub/OS/Linux`	`doc/HOWTO`	`doc/FAQ`	`doc/doc-project`
`tsx-11.mit.edu`	`/pub/linux`	`docs/HOWTO`	`docs/linux-faq`	`docs/LDP`
`sunsite.unc.edu`	`/pub/Linux`	`docs/HOWTO`	`docs/faqs`	`docs/LDP`

Files can be copied from these sites using the Internet File Transfer Protocol (FTP) [9]. Readers with access to World Wide Web (WWW) clients and the "Mosaic" browser can enjoy hypertext access to the Linux home page at the address `http://sunsite.unc.edu/mdw/linux.html`. Readers without Full Internet Access, who can only send and receive e-mail, will have to recourse to ftp-by-mail servers, as described in [9].

The first document a Linux novice should obtain is the "Installation HOWTO" by Welsh [1]. There are more than a dozen further "HOWTO" documents dealing with topics such as News, Ethernets, Printing etc. The

[2] A larger (17 inch) screen can be obtained for an additional $300. Virtual window managers, discussed later in this article, alleviate to some extent the limitations of small screens.

next most important, for the newcomer, is the "Distribution HOWTO" which describes how Linux can be obtained via FTP or from commercial vendors [2]. The cost of commercial distributions is nominal (varying from US$20–60) and they are attractive alternatives to spending hours copying diskettes via FTP. If purchasing commercial distributions, CDROM versions should be preferred whenever possible because the medium is more reliable compared with floppy disks and because much more information (usually including all documents described in this Section) is included on them.

Several Usenet newsgroups discuss Linux. The newcomer should read `comp.os.linux.announce` and `comp.os.linux.setup`. "Linux Frequently asked Questions (FAQs)", by Jackson [3] summarizes common queries on these newsgroups.

Linux supports a large number of disk controllers, video accelerator cards, Ethernet cards, monitors etc. The "Hardware HOWTO" [7] should be consulted before investing in new hardware to support Linux.

A number of books describing various aspects of Linux have been written. These can also be downloaded via FTP from the sites described above. The most important is *Linux Installation and Getting Started* by Welsh [4]. Some commercial distributors of Linux include a printed copy of this book as part of their package.

If a network is planned, the *Network Administrators Guide,* by Kirch [6] and the "Ethernet HOWTO" by Gortmaker [5] are very useful resources. *The Linux Journal* [8] is a monthly magazine and is useful for keeping up with recent developments.

The above discussion points out the large amount of information that is available about Linux. It should not discourage the newcomer: it is quite possible to install a fully operational Linux system given just a CDROM or floppy disk set and its associated documentation. A connection to the Internet, though useful, is not at all necessary.

3 The Linux Software License

Free Software Foundation, Inc., of 675 Mass Ave, Cambridge, MA 02139, USA has developed the GNU General Public License to encourage development and dissemination of free software. Most of Linux is covered by this License. The terms of this license, in summary, are as follows.

> Software can be copyrighted but free. Free Software can be distributed, a service fee can be charged if desired, but the source code should also be available if requested by the end user. The software distributed can be modified, if desired, but this fact must be made clear to the end user. This 'derived' software must continue to be freely distributable. There is no warranty for free software.

Of course a program *developed on* Linux may, if desired, be treated as a commercial product and be sold as such. However a *modification of* Linux cannot be made a commercial product. Many commercial companies distribute Linux and associated free software and charge a modest fee for this service. The purchaser of such a distribution is free to redistribute it with or without a fee. One may purchase a single CDROM or diskette set containing Linux and establish an entire laboratory with it.

4 Obtaining and Installing Linux

4.1 Obtaining Linux

Linux has traditionally been organized as a set of 'diskette series'. Thus we have the `A` series, which contains the basic system, the `D` series (compilers and development tools), the `N` series (networking) and so on. The total number of diskettes available is now about 50. However, not every user needs the entire collection. For example, if you do not want networking you need not obtain the `N` series.

There are many different Linux distributions. I shall describe the installation procedure required for the "Slackware" distribution, developed by Pat Volkerding, which is quite popular and very stable. The basic ideas described here apply to most other distributions.

There are six main methods of obtaining copies of Linux.

1. Suppose you have access to a computer connected to the Internet, and this is *not* the machine on which you wish to install Linux. This machine must have a floppy disk drive and, if it is not a DOS machine, must have software capable of writing out diskettes in DOS format[3].

[3]Most Unix workstations have the `Mtools` package for this purpose.

The File Transfer Protocol (FTP) program is to copy over diskette images from one of the Internet sites listed in Section 2.

2. To install Linux on the same machine which you are using to access the Internet, you simply need to copy the diskette images into a convenient DOS directory, and move to the installation step.

3. The most convenient method of obtaining Linux is to use the Network File System (NFS) to mount Linux directories at a remote site onto your local machine's directory structure. The remote directories appear to your system to be local sub-directories. You can now copy out the diskette sets that you want or install Linux *directly* from the NFS mounted directory.

4. Purchase diskette sets or CDROMs from a commercial distributor.

5. Borrow diskette sets or CDROMs from a friend.

6. Purchase a machine with Linux already installed from a commercial vendor.

4.2 Installing Linux

4.2.1 Hardware Requirements

At the very minimum, Linux requires an IBM-PC compatible with a 386 or higher processor, 4 MBytes of memory, and 80 MBytes of free hard disk space. Such a system is capable of running Linux and can give the novice a feel for Unix. The X Window System can run on such a system, albeit with very poor performance; compilers, such as `gcc` (GNU C), are intolerably slow. With 8 MBytes of RAM and 100 Mbytes of disk space, we get a system that is useful for program development, and can run X. The amount of RAM and hard disk space is more important than the speed of the processor. I have run Linux quite comfortably on 386/16MHz machines with 8 Mbyte RAM. A numeric co-processor is optional but will, of course, be very useful if you plan on doing heavy numeric computation.

Linux supports a large variety of monitors and graphic cards. X can run even on simple monochrome Hercules cards. In general, SVGA monitors give

better performance. A video accelerator card improves the performance of X very significantly.

A PC with a 486/DX2-66MHz processor, 8 MBytes of RAM, 200 or more Mbytes hard disk space, video accelerator card and 0.28in. pitch SVGA color monitor constitutes a fine Linux system. An additional 8 Mbytes of RAM will improve performance significantly.

Obviously, the amount of hard disk space required is proportional to the amount of software you want to load into the system. The full set of electronic/VLSI CAD tools that is available for Linux requires an additional 100 Mbytes. One warning is in order for those interested in running TEX/LATEX. The fonts required to preview or print out formatted output are generated on demand. Thus you might install the TEX/LATEX package on your system and have it up and running, only to discover a few weeks later that you have run out of disk space because the current paper you are writing requires new fonts that have filled up your disk. It is best to set aside 4–5 Mbytes disk space for future font generation.

4.2.2 Creating Boot Disks

The first step in setting up Linux is the creation of a 'boot' diskette. The Linux diskette set or CDROM that you obtained earlier will have a DOS program called `rawrite.exe` whose sole purpose is to create a boot diskette from one of several possible 'boot images' as illustrated in Figure 1. A boot diskette can be used to start up a PC and contains a minimal Linux system that runs off a floppy disk and RAM memory. This minimal system is used to prepare the hard disk and then to actually load and configure Linux.

4.2.3 Setting up the Hard Disk

Partitioning of the hard disk is illustrated in Figure 2. In this example, I assume that a hard disk with a total of 250 MBytes is available and that there is a partition of 100 Mbytes that contains DOS.

If this is a new PC, the 100 MByte DOS partition would be created at the outset. However it is more likely that an existing 250 MByte hard disk, wholly devoted to DOS, is repartitioned to allow the installation of Linux. The obvious way to do this is to back up all your personal files on floppies, use the DOS `fdisk` command to create a smaller 100 MByte partition, reinstall

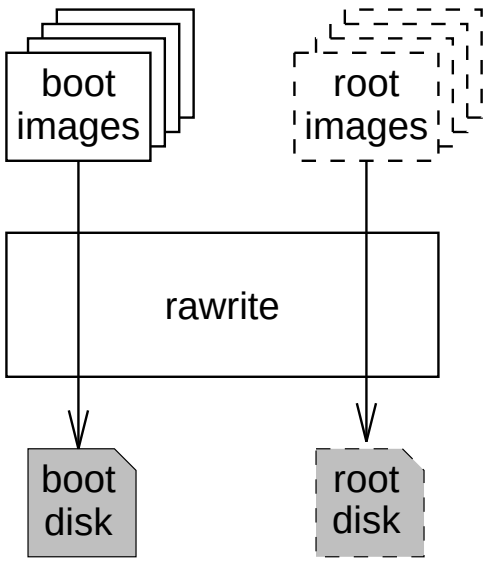

Figure 1: Creating a Linux boot diskette. A DOS program called `rawrite.exe` and a number of boot images are included in all Linux distributions. `rawrite` is used to create a PC bootable diskette, which contains a minimal Linux system that runs off a floppy disk and RAM. Some Linux distributions require an additional 'root' diskette that is inserted in the floppy drive after the 'boot' diskette has been read in. From amongst several possibilities, you select the disk image that corresponds to your specific needs (e.g. whether your machine has a CDROM, color monitor etc.).

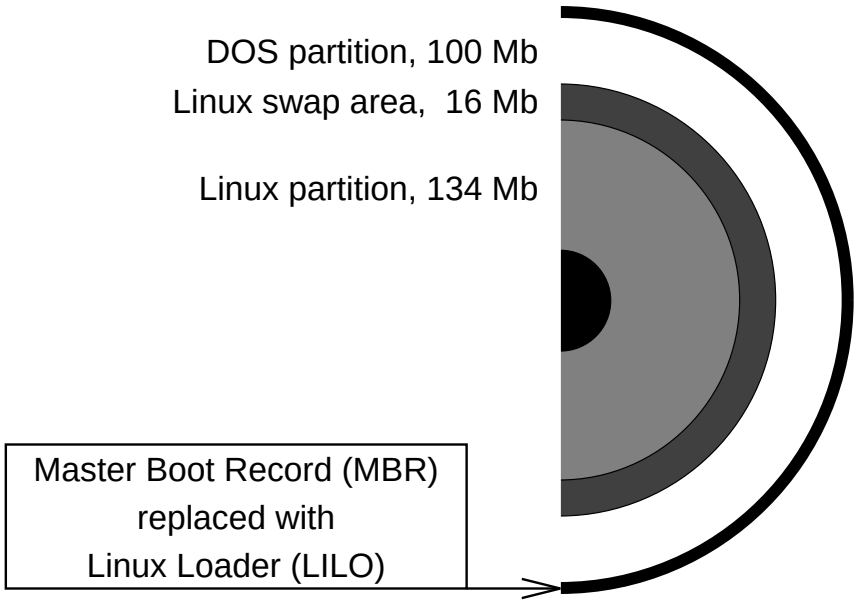

Figure 2: After starting Linux from a boot diskette, the `fdisk` program is used to partition the hard disk. In this example, a 250 MB hard disk has a 100 MB DOS partition, a 16 MB Linux swap area (to implement virtual memory) and a 134 MB partition to hold the Linux operating system. Once the complete Linux system has been installed (as described in Section 4.2.4), the Master Boot Record (MBR) of the hard disk (which contains a small program that ordinarily starts up DOS) is replaced with the Linux Loader (LILO) program to permit booting either DOS or Linux when the machine is turned on. This gives the user the flexibility of starting up his machine with either DOS or Linux.

DOS and your application software from their original floppies and your personal files from the back up floppies.

A more convenient method is to first use a DOS disk defragmentation program to compact all your files into the beginning cylinders of your hard disk and then use the `fips` program (accessible at `tsx-11.mit.edu`, see Section 2) to perform a 'non-destructive' repartitioning that leaves your original files undisturbed.

The Linux `fdisk` program is now used to create two new partitions. One is a 'swap' partition of 16 Mbytes, which is used to implement virtual memory, and the other is a 134 Mbyte partition that will hold the Linux operating system.

We must now create 'swap space' in the swap partition and make a file system in the Linux partition. These steps are analogous to the `format` operations of DOS.

4.2.4 Setting up Linux

A program called `setup` takes over the responsibility of actually installing Linux. This program prompts the user to specify the diskette sets (Section 4.1) he wishes to install and then proceeds to create appropriate directories and install relevant files.

Figure 3 illustrates this process. In part (a) Linux is being installed from a collection of diskette sets. The `setup` program will prompt the user to insert specific diskettes as installation proceeds. CDROM installation is simpler, as shown in part (b), as all diskette sets exist on the CDROM and only the selected ones are loaded. In part (c) of this figure we show Linux being installed from a DOS partition of a hard disk into a Linux partition of the same disk. The diskette sets have previously been copied over into the DOS partition, as described in Section 4.1.

Finally, in part (d) of Figure 3 we show how an NFS mounted installation proceeds. A directory containing Linux diskette sets is NFS mounted on the minimal Linux system running on the target PC. To the target PC, the Linux diskette sets appear to lie in the PC's own sub-directories, whereas in actual fact they are in a distinct machine.

Once the diskette sets have been installed, the `setup` program guides the user through the configuration of the system. This is fairly straightforward and involves things like setting up the modem, mouse, network address (if

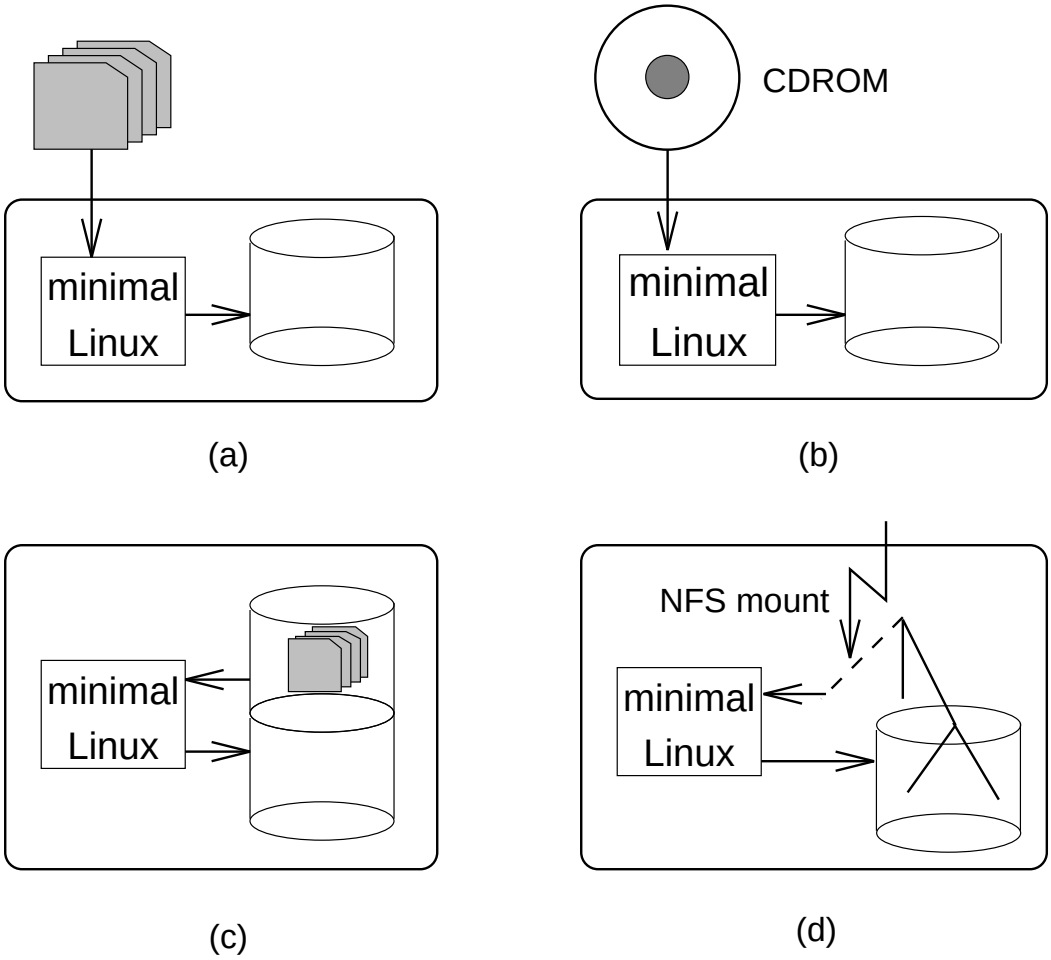

Figure 3: A PC started with a Linux boot diskette runs a minimal Linux system based on a floppy disk and RAM. This is used to 'install the complete Linux system from (a) a set of diskettes or (b) a CDROM. If the disk sets have been copied into a DOS partition of the machine, then (c) Linux can be installed from one partition to another. Linux can also be installed from an NFS mounted directory (d).

any), time zone and so forth. An important step is the installation of the Linux Loader (LILO), as described in Figure 2, which can be set up to boot Linux or a number of other operating systems.

5 Running Linux

The day-to-day running of a standalone Linux machine is not particularly difficult, especially if you are already familiar with Unix. If you are new to Unix, this will be a valuable learning experience. In common with all Unix systems, a `root` account is used for system administration, setting up user accounts, installing and removing software etc. Even if only a single individual is to use a Linux PC, it is highly recommended that routine work not be done under the `root` account, as a carelessly typed command can cause much grief: a separate user account should be created.

Shutting down Linux is not a simple matter of turning off the power switch. This is because, like most Unix systems, large amounts of disk contents are cached in RAM buffers and written out to disk at periodic intervals. Before shutting down the system, the disk must be 'synced', i.e. the RAM buffers must be flushed. This is done using the `shutdown` command. Normally `shutdown` can only be executed by `root`. However there exists a utility called `sudo` that permits authorized users to execute `shutdown` and other restricted commands, normally executed only by `root`.

It is possible to configure Linux so that the traditional DOS reboot signal 'ctrl-alt-del' causes the system to shut down gracefully. We use such an arrangement in our laboratory at Lahore. Students are normally not permitted to turn power off. At the end of the day the laboratory custodian (who does not have a user account) shuts down each machine by hitting 'ctrl-alt-del', waiting for the machine to signal that it has halted, and only then turning off power. However it is useful to inform students of this shutdown procedure, so that they can use it in an emergency, such as sudden severe weather.

The above discussion also points out the necessity of using an uninteruptible power supply (UPS), particularly in locations prone to frequent power outages (e.g. most developing nations). This is all the more important in systems used for 'production' applications. It is possible to set up Linux to shut down automatically if the power outage is so long that the UPS runs out of battery power. This requires a UPS with the capability of sending a

'power-down' signal to the computer.

When Linux is employed in a multi-user environment, particularly on a network, the problems of system administration are exactly the same as in traditional Unix systems. One or more competent system administrators are required to devote considerable effort to keep the system running. It is helpful, in this case, to have as many similar systems as possible. Ideally all machines should be the same so that software and configuration files can be automatically copied to each machine. The Network File System is useful in reducing file distribution overhead and maintaining a consistent set of software over many machines.

5.1 Linux and MSDOS

Linux can coexist peacefully with MSDOS. This means that one hard disk can hold DOS as well as Linux and either can be started up at boot time, as explained in Figure 2 above. Linux can also access DOS partitions and can *automatically* perform the conversions necessary to make text files readable by either system. This capability is required because DOS terminates each line of text with carriage-return+line feed, whereas Linux/Unix only uses line feed.

Ordinary MSDOS diskettes can be read or written by Linux using the `Mtools` set of commands. There is a project underway in the Linux community to develop DOSEMU, a DOS emulator that permits the execution of DOS programs under Linux. This project has achieved considerable success. A similar project is the MSWindows emulator, WINE. Information on these is accessible on the Internet.

6 Networking

Linux supports the TCP/IP protocol that is commonly used for intercommunication between Unix systems. Thus networks of Linux machines can be set up to provide all the functions normally associated with Unix networks. Linux machines can also be integrated into existing networks that run TCP/IP.

Setting up a network of Linux PCs is quite straightforward. Linux supports most of the common Ethernet cards. Prices for these are currently as

low as US$50. If you are interested in hooking up to an existing network in your office or campus, you will need your network administrator to assign you a network address. If you are setting up an isolated network you can use any address you wish. However it is preferable to get an address properly allocated to your site, because you may eventually want to connect to the Internet.

The `setup` program guides you in setting up network addresses, node names, etc. You can activate the Network File System (NFS) on your network to permit several machines to remotely mount directories from a specific machine. This saves disk space and administration overhead and permits students to work on a central copy of their home directories no matter where they login. Every machine in a network can 'export' its directories in this fashion, permitting very flexible operation.

7 The X Window System

X is the de-facto standard windowing system for Unix based workstations. It was developed by a consortium of companies centered around MIT. XFree86 is a port of the X Window System for the 386/486 architecture which has been developed by a team of volunteers and is included in most Linux distributions.

Many popular monitors and video cards are supported by XFree86. However configuring XFree86 to support your particular monitor and video display card is sometimes a challenging task. Configuration files for several common monitors and video cards are included in the Linux documentation or accessible on the Internet. Video *accelerator* cards, which speed up operations when running X, are also supported.

The `ConfigXF86` program helps you set up the configuration files that are required by XFree86 to properly drive your video card and monitor. Care must be taken in this process as, in theory, an improper configuration can physically harm your monitor. In practise, this rarely happens, especially with current equipment, which has builtin safeguards.

To circumvent the limitations of small display screens when running X, one can use a 'virtual' window manager, several of which are available for Linux. These give the user a desktop that is typically 3×3 times the size of his screen. Any one of these 9 'virtual' screens can be shown on the 'real' video screen by clicking on appropriate parts of a virtual window manager icon.

8 Some Interesting Applications

Many interesting public domain or copyright–but–free applications are available for Linux. These include sophisticated Computer Aided Design tools for electronic circuit design, plotting packages, signal processing programs, finite element analysis tools, etc. I shall discuss three applications I have found to be especially useful.

The **Parallel Virtual Machine (PVM)**[12] is a system for implementing a parallel computing system on a network of workstations. The low cost and high power of current 486/Pentium PCs makes PVM an interesting application for a network of Linux machines. The bottleneck in such a system is the single Ethernet that has to carry all interprocessor communications. However the low cost of Ethernet cards for PCs makes a system with multiple Ethernets very viable. The Beowulf Project [13] at CESDIS, NASA Goddard Space Flight Center, is currently developing a system that uses multiple Ethernets to implement powerful parallel machines based on low-cost Linux PCs.

The **Ghostscript** package (developed by Alladin Enterprises and distributed by the Free Software Foundation) permits the viewing or printing of PostScript documents. One of its most interesting features is the ability to transform PostScript into Hewlett-Packard or compatible laser printer command language. This permits you to print out PostScript documents on an ordinary laser printer (provided it has at least 1.5 MByte internal memory). Such printers are about half the cost of printers capable of interpreting PostScript internally.

XMX (X-Protocol Multiplexor) was developed by John Bazik at Brown University's Computer Science Department to implement an electronic classroom on a network of workstations running X. It permits one or more windows open on the instructor's machine to be displayed on all student ma-

chines.

XMX can be used to display slides, diagrams and algorithm animations. It can also be used as a tutoring aid: the instructor opens his window on everyone's screen to illustrate, say, a Unix command. Each student opens a window on his personal machine to try out the command himself. I have installed XMX on our Linux network at the Department of Electrical Engineering at the Engineering University in Lahore and have used it for an Operating Systems course. I have used `top`, a Linux variant of the standard `ps` (process status) command, to illustrate key concepts to the class. `top` prints periodic snapshots of all processes in the system, the amount of memory (real and virtual) in use, etc. This permits the instructor to point out running and stopped processes, the impact of virtual memory usage, the advantages of shared libraries, and so on. All this is done 'live' during the course of the lecture.

9 Conclusions

Linux is here to stay. It is a stable, powerful operating system that runs on cheap commodity hardware and can be freely distributed. It has been developed by a team of volunteers spread throughout the world and is a fine example of international collaboration through the Internet.

The processor family on which Linux has been implemented, the Intel 386/486/Pentium series and its clones, is the dominant architecture in the world today, and is likely to continue to be so for the next 5 years. The huge market for this processor and its ancillary equipment ensures that the cost of Linux systems will always be very low.

A port of Linux for the IBM PowerPC is currently underway. If this new microprocessor manages to break the stranglehold of the 486 on the world market, a PowerPC version of Linux will not be far behind.

Most PCs in the world run DOS/MSWindows and commercial or entertainment applications based on these operating systems. Linux will probably have a minor impact on this market. However, within the scientific, engineering and academic communities, where Unix is ubiquitous, Linux will be a major force to be reckoned with. Should a scheme be developed to execute existing commercial 386/486 Unix binaries under Linux, and such projects are underway, it may well become the dominant Unix-clone in the world—and

it will be completely free!

Acknowledgments

I wish to acknowledge the contributions of Linus Torvalds and the world-wide cast of thousands who have made Linux possible. I would like to thank Donald Becker for patiently explaining many details of Linux to me. Ashraf Chughtai and Ashraf Iqbal have given me generous assistance in setting up our network-cum-electronic classroom at Lahore. Comments by Tom Crockett, Piyush Mehrotra and the referees have helped to improve the manuscript.

References

[1] Matt Welsh. The Linux Installation HOWTO. 1994.

[2] Eric S. Raymond. The Linux Distribution HOWTO. 1995.

[3] Ian Jackson. Linux Frequently Asked Questions with Answers. 1995.

[4] Matt Welsh. *Linux Installation and Getting Started.* 1995.

[5] Paul Gortmaker. The Linux Ethernet HOWTO. 1995.

[6] Olaf Kirsch. *Linux Network Administrators Guide.* O'Reilly and Associates, Sebastopol, California, 1995. Also available in an electronic version.

[7] Ed Carp. The Linux Hardware Compatibility HOWTO. 1995.

[8] *The Linux Journal.* Specialized Systems Consultants, P.O. Box 55549, Seattle, Washington, 98155-0549, USA. A monthly magazine devoted to Linux.

[9] Ed Krol. *The Whole Internet User's guide and Catalog.* O'Reilly and Associates, Sebastopol, California, 1994.

[10] Matt Welsh and Lar Kaufman. *Running Linux.* O'Reilly and Associates, Sebastopol, California, 1995.

[11] Stefan Strobel and Thomas Uhl. *Linux: Unleashing the Workstation in your PC.* Springer–Verlag, Berlin–Heidelberg, 1994.

[12] Adam Beguelin, Jack Dongarra, Al Geist, Robert Mancheck, and Vaidy Sunderam. *PVM: Parallel Virtual Machine—A Users Guide and Tutorial for Network Computing.* MIT Press, 1994.

[13] T. Sterling, D. Savarese and D. J. Becker. BEOWULF: A parallel workstation for scientific computation. *Proc. 24th. International Conf. on Parallel Proc., August 1995.*